A Decision Support System for Radiation Therapy Treatment Planning

Ines Winz
Department of Engineering Science
University of Auckland
New Zealand
i.winz@auckland.ac.nz

Abstract

This paper reports on the use of an interactive decision support system to generate treatment plans for external beam radiation therapy. When evaluating treatment plans, radiation therapists are often faced with situations where overdose of healthy organs or underdose of the tumour cannot be avoided. Current treatment planning software incorporate trial-and-error methods to change treatment parameters and re-optimise in order to generate a better treatment plan. However, this process is not only inefficient, but also does not yield information on how radiation doses depend on the structure site and the planning parameters. Here a multi-criteria-based optimisation model is presented, which is used to calculate a large number of efficient treatment plans. These are stored in a database and accessed for evaluation by the decision support system Carina. The navigation among those solutions and the information that is provided to guide the user in this process are described. Of special importance is sensitivity analysis, which extracts dose dependence information for the tumour and healthy organs from the efficient treatment plans. As a result, plan quality is improved by finding advantageous trade-offs in competing treatment plans. Additionally, the common trial-and-error process is avoided and effectiveness in treatment planning is increased.

1 Introduction

Cancer is one of the most significant health problems worldwide with respect to its incidence and mortality alike. One of the main treatment forms besides surgery and chemotherapy is radiation therapy. Here, ionising radiation is used to damage the DNA and interfere with cell division and cell growth. An estimated 50% of all patients diagnosed with cancer would currently benefit from radiotherapy, either to cure the disease or to palliate symptoms.

Cancerous cells are more susceptible to radiation than healthy cells which is exploited by radiation therapy. This difference in susceptibility is called the therapeutic ratio. Treatment planning is concerned with improving the therapeutic ratio by choosing optimal intensities, beam directions, etc. Due to recent improvements in

medical imaging (e.g. magnetic resonance imaging) and radiation intensity modulation (e.g. multi-leaf collimation), an increase in applicability and treatment success of radiation therapy has been noted. However, narrow therapeutic ratios, which result either in a lethal dose deposited in the tumour leading to unacceptable damage to one or more healthy structures or, conversely, in ineffective treatment by avoiding any damage to healthy structures, are still widely observed and have to be dealt with.

The planning of treatments can be very complex. Tumours often have irregular shapes, and are surrounded by or are growing into nearby organs. Additionally, a large number of plan parameters, such as number of treatment fields and beam directions, form interdependent, non-intuitive relationships that influence the final radiation dose distribution. The optimisation of the radiation intensity delivered by pre-set or pre-optimised beams is managed by the treatment planning system (TPS). The optimised treatment plan is then evaluated by the radiation therapist and/or radiation oncologist. Often optimal solutions of mathematical models underlying optimisation are not feasible or clinically acceptable. This invariably results in a trial-and-error process where the planner changes input parameters in the search for a better optimisation output. This search may be very time-consuming, depending on the experience of the planner and the complexity of the case.

This paper reports on a different approach to treatment planning that avoids the trial-and-error process. Instead, a decision support system provides the planner with the necessary guidance in selecting the final treatment plan and making trade-off decisions between a set of pre-computed treatment plans. Based on multi-criteria optimisation, this approach has received increasing attention over the past years, as it deals more effectively with the problem of narrow therapeutic ratios.

2 Radiotherapy Treatment Planning

A treatment plan consists of the equipment configuration, resultant dose distribution in the patient, and a set of treatment instructions. Activities leading to a finalised treatment comprise: patient positioning and immobilisation, imaging, delineation of structures, plan optimisation, plan verification, treatment, quality assurance and verification.

Treatments can have either a curative or palliative intent. A curative treatment is focused on generating a tumouricidal dose, while the total dose in a palliative treatment is comparatively lower and is aimed at achieving temporary relief of symptoms. Dose is typically measured as absorbed radiation in units of $Gray\,(Gy)$. These primary objectives are realised by specifying a number of treatment goals. These include avoidance, conformity, homogeneity, and simplicity goals. Guidelines and protocols usually recommend values for each of these goals with respect to treatment site and progress of the disease. Despite this, the level of importance of each of these goals may vary and radiation therapists may use them differently.

Classical, *forward* treatment planning is conducted iteratively. The radiation therapist specifies all parameter values, after which the dose calculation software computes the dose distribution. The initial parameter values are then adapted by trial-and-error until the dose distribution is satisfactory. This planning strategy is still used in about 90% of the 5,500 cancer centres worldwide (Varian Medical

Systems: Annual Report 2002).

Recently, an *inverse* planning approach has become popular and necessary to cope with the increased complexity of intensity-modulated beams. In inverse planning, the software computes optimal beam intensities given specified requirements for the maximum and minimum dose deposit in organs and tumour. Increasing the number of requirements to account for these problems often results in an empty feasible region. This leads to a similar trial-and-error process, where dose requirements are adjusted so that a feasible solution can be produced which is also clinically acceptable.

For forward and inverse planning alike, the underlying problem is that minimising the dose to healthy tissues conflicts with generating a sufficiently high dose to the tumour. As a result, difficult decisions have to be made regarding the overdosing of organs and/or the underdosing of the tumour. These decisions will always have to balance the perceived risk of unsuccessful tumour control with the possibility of complications to healthy tissues.

Our approach is aimed at generating a well-balanced, optimal trade-off between the over- and underdose of organs and tumours, respectively. Multi-criteria optimisation is a conclusive way of dealing with the conflicting objectives of over- and underdosing healthy structures and the tumour. As a result, the trial-and-error process is abandoned in favour of guided search among pre-computed Pareto-optimal treatment plans (Figure 1). Pareto-optimality or efficiency implies that no objective can be improved without deteriorating another objective. From this follows that a treatment plan chosen for a patient should always be Pareto-optimal.

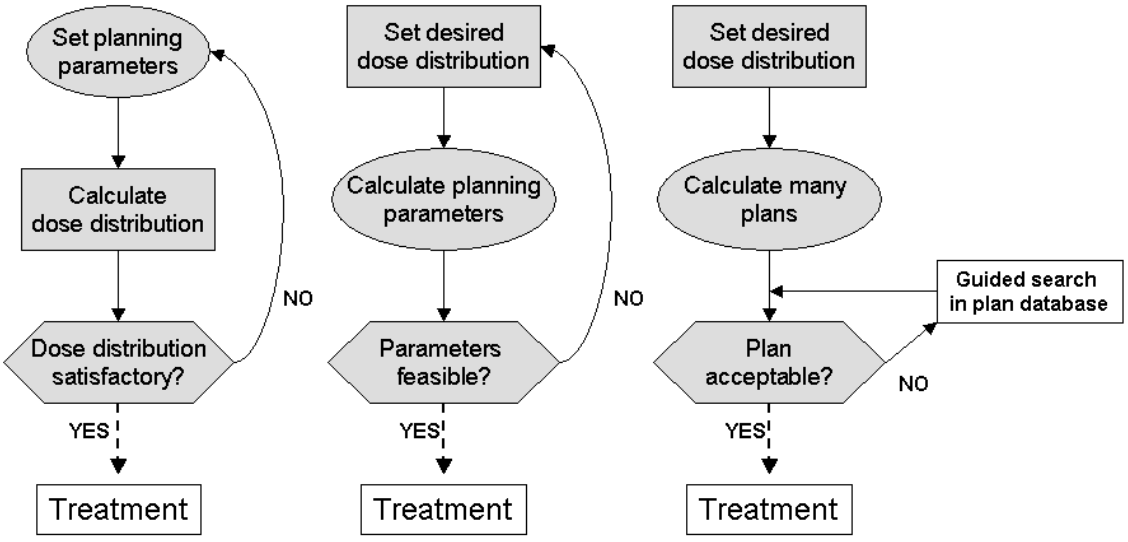

Figure 1: *Changes in the treatment planning paradigm from forward planning (left) to inverse planning (middle) to decision support-based planning (right).*

3 Decision Support-Based Treatment Planning

3.1 Current Practice

The intensity problem is the basic optimisation problem in radiation therapy. It is aimed at finding the best radiation intensities in all beam heads. A large number of different optimisation models (LP, MIP, NLP) have been suggested to optimally

solve this problem. Currently widely implemented in TPS is an NLP model that minimises the weighted and squared differences between actual and prescribed dose in all structures. It uses weights to account for the importance of each structure. However, weights represent an artificial rather than a clinical concept of structure importance, and hence they are hard to assess and determine. Furthermore, even slight changes in weight factors were shown to radically change the optimisation result. This results in mathematically optimal treatment plans that are clinically irrelevant.

3.2 Multi-criteria problem formulation

The multi-criteria formulation of the intensity problem is a direct consequence of the fact that radiation therapy treatment planning deals with conflicting objectives (Ehrgott and Burjony 2001; Küfer et al. 2003). The mathematical model is based on the discretisation of the body into volume elements (voxels) and the radiation beam into beam elements (bixels). This discretisation leads to large-scale constraint matrices.

The following multi-criteria linear programme (Hamacher and Küfer 2002) was implemented into our decision support system CARINA:

Indices

k = structure index ($k = 1$ tumour, $k = 2, \ldots, S$ organs).

Parameters

T = vector, deviation between prescribed and actual dose in all structures,
D = vector, dose deposited in all structures,
L = vector, lower prescribed dose bounds for all structures,
U = vector, upper prescribed dose bounds for all structures,
P = matrix of dose deposited in all voxels from unit intensity in x,
e = vector of ones of suitable length,
DVH_k = the dose-volume histogram function for structure k,
$\widehat{U}$ = vector of overdose accepted in a dose-volume restriction for all structures,
Υ = vector of volume percentage that may receive overdose in a dose-volume restriction for all structures.

Decision variables

x = vector of radiation intensities in all bixels.

Model MCLP

$$\textit{minimise} \qquad T(x) \;=\; \{T_1(x), \ldots, T_S(x)\}\,, \tag{1}$$

where the deviation $T_k(x)$ is defined as:

$$\begin{aligned} T_1(x) &= \|(L_1 e - D_1)_+\|_\infty \\ T_k(x) &= \|(D_k - U_k e)_+\|_\infty \qquad\qquad \text{for } k = 2, \ldots, S. \end{aligned}$$

The constraint set is as follows:

$$
\begin{aligned}
D_1 &= P_1 x \geq (L_1 - T_1)e && (2)\\
D_k &= P_k x \leq (U_k + T_k)e \qquad \text{for } k = 2, \ldots, S && (3)\\
DVH_1\left(\widehat{U}_1\right) &\geq \Upsilon_1 && (4)\\
DVH_k\left(\widehat{U}_k\right) &\leq \Upsilon_k \qquad \text{for } k = 2, \ldots, S && (5)\\
(1-s)P_{ref}\, x &\leq P_1 x \leq (1+s)P_{ref}\, x && (6)\\
x &\geq 0. && (7)
\end{aligned}
$$

Explanation

(1) In the multi-criteria formulation, deviations T from the prescribed dose are considered separately for each structure. The objective is to minimise the maximum deviation from specified bounds in all structures.
(2) A lower bound limits the dose deposited in the tumour.
(3) Upper bounds limit the dose deposited in the organs. Note that constraints (2) and (3) are easily implemented, but assume that all organs have the same dose-response relationship.
(4) Minimum dose-volume constraint: "At least $\Upsilon_1\%$ of the tumour should exceed $\widehat{U}_1\, Gy$": In practice, it is often acceptable to deliver doses greater than the bound, as long as the partial fraction of the organ receiving the excess dose is small. These dose-volume relationships (DVR) are widely used by radiation therapists. DVR are most efficiently modelled and implemented by approximation, as otherwise a mixed-integer formulation is necessary. This greatly increases the problem size and makes it harder to solve.
(5) Maximum dose-volume constraint: "No more than $\Upsilon_k\%$ of structure k should exceed $\widehat{U}_k\, Gy$".
(6) The dose deposited in the tumour is compared to, and allowed a certain deviation s from, the dose in a reference voxel P_{ref} of the dose matrix P. This constraint enforces tumour dose homogeneity. Very inhomogeneous dose distributions can lead to adverse medical conditions.
(7) Finally, non-negativity constraints ensure that no negative beam energy is computed. By definition, T is always non-negative.

It has to be noted, that due to the problem of narrow therapeutic ratios the ideal solution $T = 0$ usually does not exist. As a result of the multi-criteria model, there will be no one best solution, but a number of Pareto-optimal or efficient solutions. These are characterised by the fact that an improvement in one criterion invariably leads to a deterioration in one or several other criteria. This is the trade-off the planner has to take into account when looking for a plan that is most beneficial for the patient.

Pareto-optimal solutions of MCLP are found as follows. For each point in a grid, a linear programme is solved, where the deviation values for all organs are pre-set and the tumour deviation value is minimised. Constraints are as specified by the radiation therapist. On average, 2000 treatment plans are computed. The calculation of Pareto-optimal treatment plans need not be supervised, and thus can be done overnight.

3.3 Grid Definition

To enable an efficient navigation and calculation process, the pre-computed Pareto-optimal solutions are arranged in a grid. It is not necessary for the grid to cover the whole solution space, as some solutions are clearly undesirable to the radiation therapist and thus will never be selected. For example, there are always solutions where $T_k = 0$ for some k. Consequently, these solutions have a high deviation value T_k in some other k. As a result, the solution space can be divided into parts of interest to the radiation oncologist and the remaining space. Only the former should be covered by the grid of efficient solutions. The range of values for each structure should be equal or similar in this grid, so that no structure is preferred over another and no bias occurs. From this follows that the starting solution will have equal or similar deviation values in each structure, and thus will be situated in the centre of the grid of Pareto-optimal solutions. This solution will be referred to as the balanced solution. Mathematically it represents the solution whose criteria have equal deviation from the ideal point $T = \{0, \ldots, 0\}$, i.e. the theoretical point where each criterion reaches its optimum value. At this point all dose bounds are satisfied without any under- or overdosing.

The grid will furthermore be homogeneous, in that the distances between all neighbouring solutions are equal. One exception to this is the division of the grid into an inner and outer grid. The *inner grid* will surround the balanced solution, and itself will be surrounded by the outer grid. The points in the *outer grid* will be less densely spaced than in the inner grid (Figure 2). This is derived from the fact that the inner grid is of higher clinical importance, as the trade-offs between structures are less extreme.

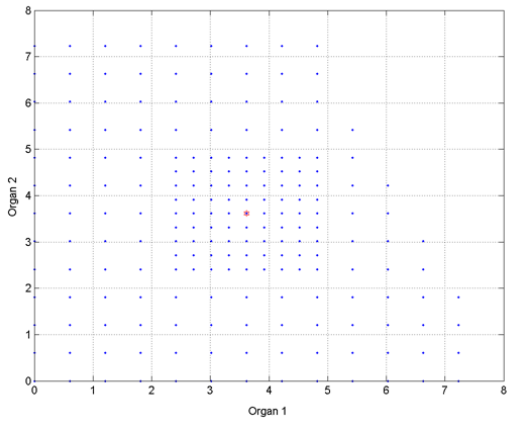

Figure 2: *The grid of Pareto-optimal solutions. The grid is not complete as the boundary of the solution space has been reached.*

3.4 Navigation

During navigation, plan evaluation and comparison take place, and the radiation therapist decides on the final treatment plan that best fulfils individual treatment goals. This requires close interaction between the software and the radiation therapist. During the evaluation of the balanced solution the radiation therapist decides which treatment goals are met or which deviation value has to be improved (i.e. decreased). For this improvement another criterion value has to be traded-off, i.e.

the radiation therapist determines which structure's deviation will be deteriorated (i.e. increased). Given these requirements for improvement and deterioration, a free search or a fine search can be initialised. In the *free search*, the radiation therapist specifies and inputs the exact values for improvement and deterioration for each structure. Another option is to fix a number of structures at their current deviation levels. CARINA writes an SQL string according to these inputs, queries the database, and returns the plan which best fulfils these requirements. In this process, the deviation between a specified value and the actual value in the database will be minimised. The free search is cumbersome when simply a neighbouring solution is sought, i.e. a slight increase in one structure traded-off against a slight decrease in another structure. In this case, a *fine search* is more appropriate. The only inputs necessary are which structure to improve and which to deteriorate. CARINA will return a unique neighbouring solution if one exists. Otherwise it informs the planner that his requests cannot be met and resets the last user inputs. A third option is an *exact search*, where the plan identifier is input and the corresponding plan is immediately displayed. This is useful when solutions are revisited for comparison purposes or after sensitivity analysis (see below).

Each iteration of a free or fine search constitutes one search process. Navigation may consist of several of these search processes until a satisfying treatment plan is found. Searches will be very fast and query times instantaneous due to the query optimisation engine in the database management system.

During each successive search process CARINA provides further information based on the plan data available in the database. It provides the user with the number of available solutions in the database, which is bound to decrease as only a limited number of plans can satisfy set requirements. CARINA also outputs available value ranges for each structure, as they also change with specified requirements. However, most important is the use of *sensitivity analysis* to retrieve information on how the dose to one structure is dependent on the dose to another structure. For example, sensitivity analysis can yield information on how much a critical structure could be spared if a tumour dose reduction is accepted. This information is not provided by current TPS, because multi-criteria methods are necessary to obtain and exploit it. Additionally, sensitivity analysis has the power to reveal advantageous trade-offs where the total improvement greatly exceeds the total deterioration when comparing two rival treatment plans (Winz 2004).

3.5 Example Treatment Planning Session

This example treatment planning session demonstrates the use of sensitivity analysis and compares it with the fine search. In this example, the tumour is situated in close proximity to three organs: spinal cord, left and right kidney. The treatment parameters specified before optimisation are: 10 treatment fields, 15 degrees equispaced beam directions, 16 MV beam energy. No dose-volume or inhomogeneity restriction was specified. The grid size is set to 9 points for each structure in the inner grid and 6 points in the outer grid. The structure bounds are: 80 Gy tumour, 33 Gy left kidney, 33 Gy right kidney, and 25 Gy spinal cord. Total plan calculation time was 3 hours and 17 minutes. 1634 solutions were calculated of which 1581 were stored as efficient in the database.

The radiation therapist's task is to trade-off over- and underdose to organs and

tumour, respectively, with regard to the treatment goals. Figure 3 shows the dose distribution diagram and dose-volume histogram for the balanced solution. The deviation values in all structures are equal at $T = 2.97$ in all structures. Doses are already close to the bounds, which is a direct consequence of the large number of treatment fields. However, there is still room for navigation. Treatment goals are such that the dose to the radio-sensitive spinal cord ($k = 4$) should be decreased further, while accepting increased irradiation of the left kidney ($k = 2$). The right kidney ($k = 3$) should be spared as well, so that the main radiation burden is on the left kidney. If possible the dose deposited in the tumour should not be more than $3\,Gy$ below the lower bound, i.e. $T_1 \leq 3$.

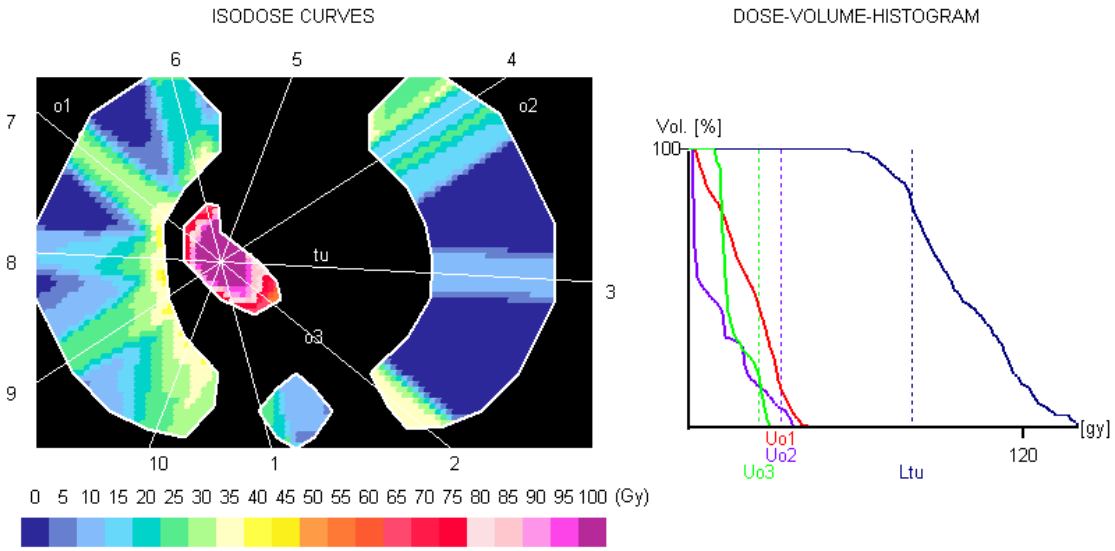

Figure 3: *The isodose plot (left) and dose-volume histogram (right) for the balanced solution. The objective function values are* $T = \{2.97, 2.97, 2.97, 2.97\}$.

With this in mind, an improved plan would have the right kidney and the spinal cord irradiated at their set limits, i.e. $T_3 = 0$ and $T_4 = 0$. In order to obtain information on dose dependence between the left kidney and the tumour, a sensitivity analysis is performed which fixes deviation levels T_3 and T_4 at $0\,Gy$. The result is given in Table 1.

planID	T_1	T_2	ΔT_1	ΔT_2	ΔT_2-ΔT_1
631	2.69	3.57	0.28	.60	.32
727	1.59	4.17	1.38	1.19	−.19
823	0.48	4.76	2.49	1.79	−.70

Table 1: *Information extracted on the dose dependence between the tumour* (T_1) *and the left kidney* (T_2).

Opting for solution 631 instead of the balanced solution represents an advantageous trade-off. The total improvement is $6.22\,Gy$, which represents the combined improvements of the right kidney ($2.97\,Gy$), the spinal cord ($2.97\,Gy$), and the tumour ($0.28\,Gy$). This has been traded-off against a total deterioration of only $0.6\,Gy$ (in the left kidney only). Solution 631 is made the new plan of choice (see Figure 4 for the dose distribution and dose-volume histogram).

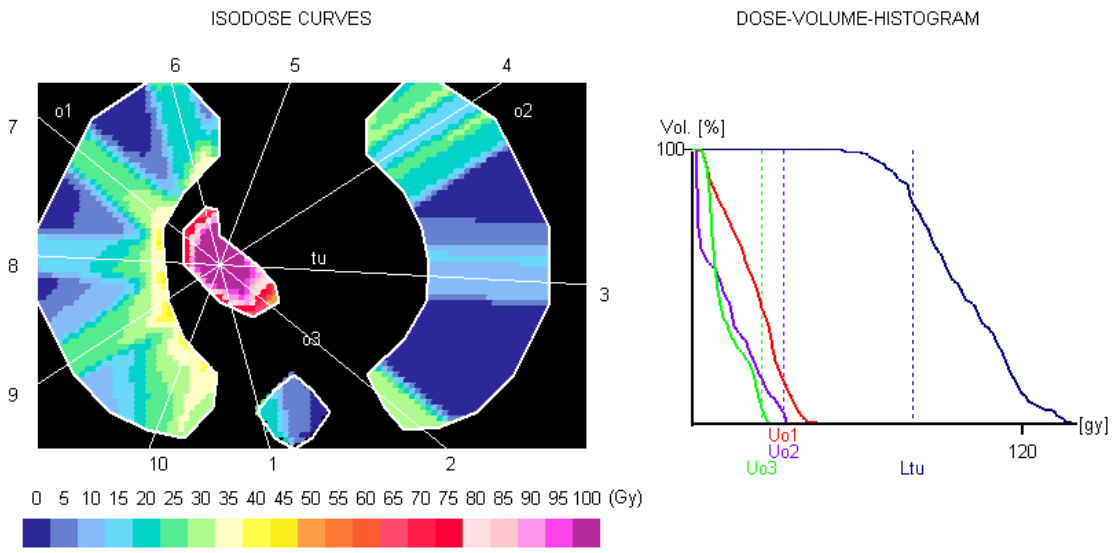

Figure 4: *The isodose plot (left) and dose-volume histogram (right) for solution 631. The objective function values are* $T = \{2.69, 3.57, 0, 0\}$.

Note that a free search with $T_1 < 3$ and $T_3 = T_4 = 0$ would have given the same result, albeit without dose dependence information.

The next step is to search within the immediate neighbourhood of plan 631 in order to achieve a better deviation value for the left kidney. A fine tuning search is done with $T_1 < 3$, $T_2 < 3.57$, $T_3 = 0$ and $T_4 > 0$. Querying the database brings forward that such a plan does not exist. Repeating the query with $T_4 = 0$ and $T_3 > 0$ shows that T_3 must be increased to at least $4.76\,Gy$. This is not acceptable. The remaining options are to either increase the tumour deviation, keep the T_2 value at $3.57\,Gy$, or decrease both T_3 and T_4. The last approach was tested by sensitivity analysis. The result is given in Table 2.

planID	T_1	T_3	T_4	ΔT_1
1213	2.998	2.68	4.17	−.306
1217	2.997	2.97	2.68	−.304
588	2.991	3.57	1.19	−.298
1224	2.983	3.27	2.08	−.290
1218	2.977	2.97	2.97	−.285

Table 2: *Sensitivity analysis for an improvement in the left kidney and a deterioration in the right kidney, the spinal cord, and the tumour.*

Unfortunately, the doses to the right kidney and spinal cord would have to be increased quite dramatically in order to decrease the dose to the left kidney even slightly. This is not acceptable. Hence, plan 631 is chosen as the final treatment plan for this patient. The corresponding beam intensities are stored in the database and can be retrieved effortlessly.

4 Conclusions and Directions for Future Research

Based on the reasonable assumption that oncology personnel want as much influence on the final treatment plan as possible, CARINA's navigation process accomplishes this using versatile options for navigation searches. Moreover, it provides information to support all search processes. Sensitivity analysis has emerged as an extremely useful option to find treatment plans having advantageously traded-off deviation values.

The decision making task performed by the radiation therapist and oncologist is supported, not replaced, by a decision support system. As a result of avoiding trial-and-error and re-optimisation, planning times are drastically shortened. In addition, plan quality is improved by finding and exploiting advantageous trade-offs.

Radiation therapy treatment planning is a very complex process. It is important to direct attention to not only the intensity problem itself, but also to factors influencing it. Consequently, radiation therapists should be supported in their choice of initial parameters. Here, beam direction optimisation is most critical, as the influence of beam directions on the resulting dose distribution is substantial and the choice of beam directions often non-intuitive (Ehrgott and Johnston 2003).

The support through sensitivity analysis should be expanded. The idea is to initialise a search program that extracts advantageous trade-offs from the database of pre-computed treatment plans. As a result, CARINA could propose treatment plans based on such a search outcome.

Acknowledgements

This research was supported by The University of Auckland's Vice Chancellor's Development Fund, Project 23102.

References

Ehrgott, M., and M. Burjony. 2001. "Radiation therapy planning by multicriteria optimisation." *Proceedings of the 36th Annual Conference of the Operational Society of New Zealand*. 244–253.

Ehrgott, M., and R. Johnston. 2003. "Optimisation of Beam Directions in Intensity Modulated Radiation Therapy Planning." *OR Spectrum* 25 (2): 251–264.

Hamacher, H.W., and K.-H. Küfer. 2002. "Inverse radiation therapy planning: a multiple objective optimization approach." *Discrete Applied Mathematics* 118 (1-2): 145–161.

Küfer, K.-H., A. Scherrer, M. Monz, F. Alonso, H. Trinkaus, T. Bortfeld, and C. Thieke. 2003. "Intensity-modulated radiotherapy - a large scale multi-criteria programming problem." *OR Spectrum* 25:223–249.

Varian Medical Systems: Annual Report. 2002. IMRT - Targeting Cancer. Report available online. `http://www.varian.com/comp/2002/var_02_imrt.pdf`.

Winz, I. 2004. "A decision support system for radiotherapy treatment planning." Master's thesis, Department of Engineering Science, School of Engineering, University of Auckland, New Zealand.